Learning Auctions with Robust Incentive Guarantees

{ Jacob Abernethy, Rachel Cummings, Bhuvesh Kumar, Jamie Morgenstern} @ Georgia Tech, Samuel Taggart @ Oberlin College

Introduction

We study the problem of designing revenue maximizing repeated auctions.

Traditional auction design requires knowledge of a prior distribution over buyer valuations and constructs an auction as a function of that prior to maximize some objective. For the revenue objective, Myerson's auction (Myerson81) characterizes the Bayesian-optimal auction.

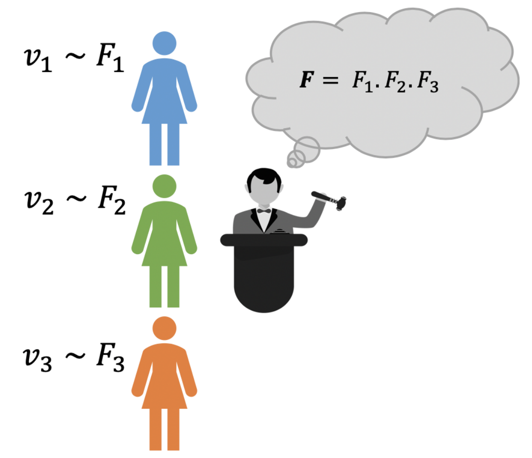

Known Priors (Eg Myerson81)

Existing work has largely focused on the computational and sample complexity of the task ignoring the possibility that bidders we learn from show up in the learned auction.

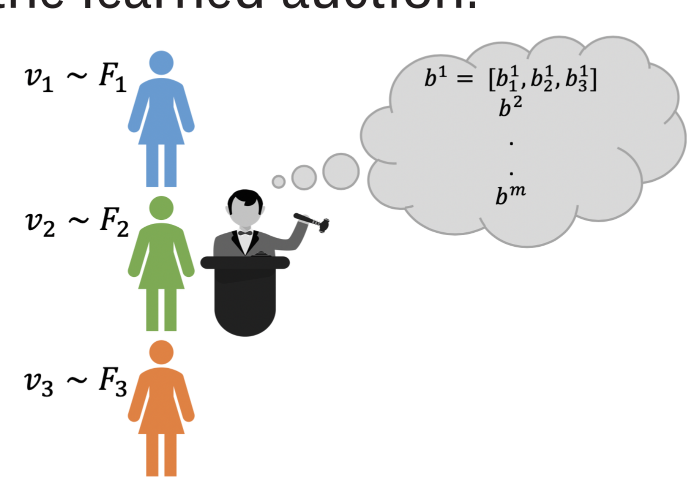

Sample Complexity results (Eg CR15, DHP15, MR16)

Our work focuses on an iterative setting

1. The seller does not know the prior and learns from the samples
2. Strategic bidders show up in multiple rounds

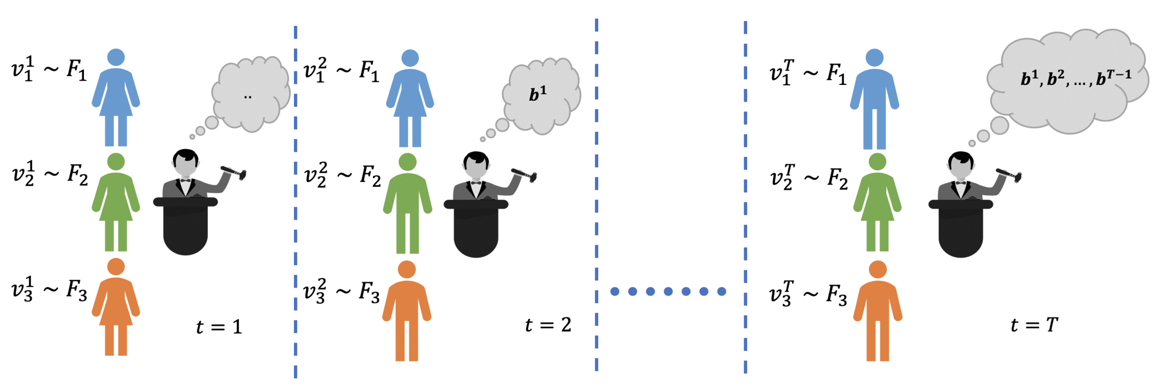

Multiple rounds, strategic bidders show up multiple times (Eg LHW18, **ACKMT19**)

Assumption

No bidder participates in more than k rounds of the T-round auction.

Setting

We consider a T-round auction, where a seller sells a supply of J identical items each round. n different populations of bidders & one bidder from each population i participates in each round t with value $v_{i;t}$ for winning an item. Each sample from population i is distributed according to a distribution D_i.

Key Ideas

1. **Differential Privacy:** Control the impact of any individual buyer's strategy on her future utility
2. **Similar revenue on similar distributions:** We show that if $k\mathbf{D} \quad \mathbf{D}^0 k_1$ then $j\text{Rev}(M; \mathbf{D})$ $\text{Rev}(M; \mathbf{D}^0)j \quad O(n^2 \quad)$

DP distribution estimation

We use a 2 fold tree based aggregation protocol (LHW18) with appropriate noise to estimate the value distributions of the bidders in a (; =T) differentially private manner.

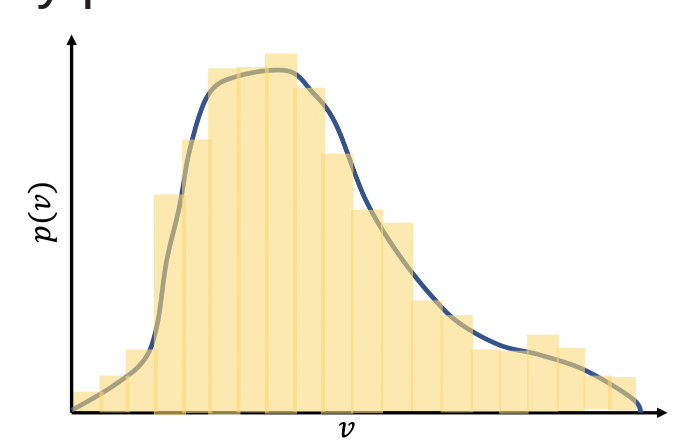

Differential private empirical distribution estimation

Exponential Mechanism

Exponential Mechanism (MT07) is an DP algorithm to sample allocation x as $P[x] \propto \exp \frac{Q(\hat{}(\mathbf{b}); x)}{}$ where the quality function $Q(\hat{}(\mathbf{b}); x)$ we use is $\sum_{i=0}^{n} \hat{}_i(b_i)x_i$

Black Box Payments

We use a black box payment rule from APTT04 to decide payments which are truthful in expectation and bidder i's payment depends on her own bid, and the differentially private estimates of the distributions but does not depend on other's bids.

Utility-Approximate BIC Online Auction

Algorithm 1: Utility-Approx BIC Online Auction

Parameters: discretization , privacy , upper bound h, num. of rounds T
Initialize: $H^0_{i;0}$ Uni$(0; h)$ for $i = 1; \quad ; n$
for $t = 1; \quad ; T$ **do**
  Receive bid profile $\mathbf{b}_t$ rounded down to multiple of
  Run Myerson with $\mathbf{H}^0{}_{t \ 1}$ as prior and $\mathbf{b}_t$ as bid for allocations / payments.
  for $i = 1; : : : ; n$ **do**
    Update $H^0_{i;t}$ via two-fold tree aggregation
  end
end

Results:

Theorem 1. *Algorithm 1 is $kh \quad 2 + \frac{1}{T}$ -utility approximate BIC, i.e for each bidder i, if everyone else bids truthfully then truthful bidding is an approximately utility-maximizing behavior for bidder i.*

Theorem 2. *With probability at least 1 , the average expected revenue obtained by Algorithm 1 for T rounds OPT J $4hn^2O \sqrt{\frac{\log(\frac{nT}{})}{T}} + \frac{1}{T}$ for regular distributions D and < 1.*

Bid-Approximate BIC Online Auction

Algorithm 2: Bid-approx BIC online Auction

Parameters: discretization , privacy , upper bound h, num. of rounds T
Initialize: $H^0_{i;0}$ Uni$(0; h)$ for $i = 1; \quad ; n$
for $t = 1; \quad ; T$ **do**
  Receive vector of bids $\mathbf{b}_t = (b_{1;t}; : : : ; b_{n;t})$ rounded down to multiple of
  With probability $_t$ select a random allocation and reserve price
  else
    for $i = 1; : : : ; n$ **do**
      Use $H^0_{i;t \ 1}$ to calculate
      $_{i;t}(b_{i;t}) = b_{1;t} \quad \frac{1 \quad H^0_{i;t \ 1}(b_{1;t})}{h^0_{i;t}(b_{1;t})}$
    end
    Use exponential mechanism to select allocation
    Use *Black box payments* to calculate payments $p_t(\mathbf{b}_t)$.
  end
  for $i = 1; : : : ; n$ **do**
    Update $H^0_{i;t}$ via two-fold tree aggregation
  end
end

Results:

Theorem 3. *Algorithm 2 is is $_t$-bid approximate BIC in round t for $_t = c + O(t^{1=4})$ for small constants c^0, i.e there exists a Bayesian Nash Equilibrium where in every round t, every bidder bids within $_t$ of their truthful value.*

Theorem 4. *With probability at least 1 , the average expected revenue obtained by Algorithm 2 for T rounds OPT c^0 $O(T^{1=4})$ for small constant c.*

Link to paper